arXiv:1605.00071v1 [math.OC] 30 Apr 2016

The Homotopy Method Revisited: Computing Solution Paths of ℓ_1-Regularized Problems

Björn Bringmann, Daniel Cremers, Felix Krahmer, Michael Möller

May 3, 2016

Abstract

ℓ_1-regularized linear inverse problems are frequently used in signal processing, image analysis, and statistics. The correct choice of the regularization parameter $t \in \mathbb{R}_{\geq 0}$ is a delicate issue. Instead of solving the variational problem for a fixed parameter, the idea of the homotopy method is to compute a complete solution path $u(t)$ as a function of t. In a celebrated paper by Osborne, Presnell, and Turlach, it has been shown that the computational cost of this approach is often comparable to the cost of solving the corresponding least squares problem. Their analysis relies on the one-at-a-time condition, which requires that different indices enter or leave the support of the solution at distinct regularization parameters. In this paper, we introduce a generalized homotopy algorithm based on a nonnegative least squares problem, which does not require such a condition, and prove its termination after finitely many steps. At every point of the path, we give a full characterization of all possible directions. To illustrate our results, we discuss examples in which the standard homotopy method either fails or becomes infeasible. To the best of our knowledge, our algorithm is the first to provably compute a full solution path for an arbitrary combination of an input matrix and a data vector.

1 Introduction

In recent years, sparsity promoting regularizations for inverse problems played an important role in many fields such as image and signal analysis [11, 3] and statistics [12]. The typical setup is that one tries to recover an unknown signal $\hat{u} \in \mathbb{R}^N$ from linear measurements $f = A\hat{u} + \eta \in \mathbb{R}^m$ under the model assumption that $\hat{u}$ is approximately sparse, i.e., has only few significant coefficients. Here η is typically small and represents measurement noise. A common approach is to minimize the energy functional

$$u \mapsto E_t(u) := \frac{1}{2}\|Au - f\|_2^2 + t\|u\|_1 \ ,$$

where $A \in \mathbb{R}^{m\times N}$ is a matrix, $f \in \mathbb{R}^m$ the data vector, and $t > 0$ is a regularization parameter. In the statistics literature, this method is called the Lasso [12], while the inverse problems literature mostly refers to it as ℓ_1-regularization. The choice of the regularization parameter t often proves to be difficult. While choosing a small t yields unnecessarily noisy reconstructions, choosing a large t diminishes the features of the original signal $\hat{u}$. An approach to this problem is to compute a minimizer $u(t)$ for every $t > 0$, and subsequently choose a suitable regularization parameter, for instance by visual inspection of this family of solutions.

A popular algorithm to compute the full solution path is the so-called homotopy method. The homotopy method is based on the observation that there always exists a piecewise linear and continuous solution path $t \mapsto u(t)$, see Figure 1 for an example. The main idea of the homotopy

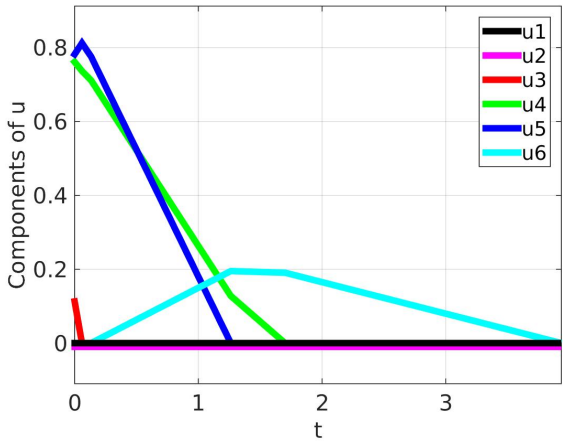

Figure 1: We plot the different components of $u(t)$ corresponding to $A \in \mathbb{R}^{3\times 6}$ and $f \in \mathbb{R}^3$ with i.i.d. standard Gaussian entries A_{ij} and f_i. The solution path $u(t)$ is piecewise linear.

method is to start at a large parameter t^0, so that the unique solution is $u(t^0) = 0$, and then follow the solution path in the direction of decreasing t. At every kink of the solution path, the classical homotopy method [10, 6, 5] computes a new direction by solving a linear system. As it turns out, the computational cost of the classical homotopy method is often comparable to the cost of solving a single minimization problem $u(t) \in \operatorname{argmin}_u E_t(u)$ or solving the least squares problem $u = A^\dagger f$. For this reason the homotopy method has proven to efficiently compute reconstructions when the noise level is unknown - provided the output is really a solution path. This has been shown to be the case given a so-called *one-at-a-time* condition [10, 6, 13], cf. Definition 12. Loosely speaking, this conditions requires that at every kink only one index joins or leaves the support of $u(t)$. The first works [10, 6] additionally required the uniqueness of the solution path $t \mapsto u(t)$, i.e., they required that $u(t)$ is the only minimizer of E_t for every $t > 0$. In [13], the homotopy method is extended to the case of non-uniqueness; again the analysis implicitly assumes the one-at-a-time condition (see Section 5.1 for a detailed discussion).

The one-at-a-time condition is known to hold in various scenarios. For example, empirical observations indicate that it is true with high probability for input $A \in \mathbb{R}^{m\times N}$ and $f \in \mathbb{R}^m$ drawn from independent continuous probability distributions. Also uniqueness of the solution path holds in many cases. Necessary and sufficient conditions for the uniqueness of minimizers have been established in [7, 14, 15]. In [13], it has been shown that, if the columns of A are independent and drawn from continuous probability distributions, the minimizer $u(t)$ is almost surely unique for every $f \in \mathbb{R}^m$ and $t \in \mathbb{R}_{>0}$.
However, both the one-at-a-time condition and uniqueness are known to be violated in certain cases [9, 13]. For instance, when the entries of A and $\hat{u}$ are chosen as independent random signs and the measurements are exact, i.e., $\eta = 0$, the one-at-a-time condition is regularly violated. In such cases it has been observed that standard homotopy implementations can fail to find a solution path [9]. If in addition uniqueness is violated, even the finite termination property of the homotopy method may no longer hold, see Proposition 9.

In this paper, we propose a generalized homotopy method, which addresses these issues. In contrast to the classical homotopy method [10, 6], which solves a linear system at each kink, the

generalized homotopy method solves a nonnegative least squares problem. The main result of this paper, Theorem 10, shows that this new algorithm always computes a full solution path in finitely many steps, even without a one-at-a-time assumption. Along the way, we give a full characterization of all directions which linearly extend a given partial solution path, see Theorem 4. Our characterization is of interest even under the one-at-a-time condition, since it provides a unified treatment of both hitting and leaving indices (cf. [6]). We also show that, under the assumptions of [10, 6], the generalized homotopy method and the standard homotopy method coincide.

1.1 Outline

In Section 2 we set up our notation and recall some basic facts commonly used in the sparse recovery literature. The set of all possible directions (cf. Definition 3) is characterized in Section 3. In Section 4 we propose the generalized homotopy method, and prove that it always computes a solution path. In Section 5 we compare the generalized homotopy method with the standard homotopy method and the adaptive inverse scale space method [2].

2 Notation and Background

For $A \in \mathbb{R}^{m\times N}$ we will denote the i^{th} column of A by $A_i \in \mathbb{R}^N$. Similary, for a subset $\mathcal{S} \subseteq [N]$, $A_{\mathcal{S}} \in \mathbb{R}^{m\times|\mathcal{S}|}$ is the submatrix of A with columns indexed by $\mathcal{S}$. Furthermore, with a slight abuse of notations, we write $A_{\mathcal{S}}^T = (A_{\mathcal{S}})^T$. The pseudoinverse of A is denoted by $A^\dagger$.
For $t \in \mathbb{R}_{\geq 0}$ and $u \in \mathbb{R}^N$, the equicorrelation set $\mathcal{E}(t,u)$ is defined as

$$\mathcal{E}(t,u) := \{i \in [N] \colon |A_i^T(Au - f)| = t\} \ . \tag{1}$$

Indeed, for least squares solutions $u_{\text{LS}} \in \operatorname{argmin}_u \|Au - f\|_2^2$, we have that $A^TAu_{\text{LS}} = A^Tf$. Even though this equation is no longer true for solutions of the ℓ_1-regularized problem, it turns out to be useful to distinguish indices $i \in [N]$ according to the magnitude of $A_i^T(Au - f)$. The active set $\mathcal{A}(u)$ is the support of u, i.e.,

$$\mathcal{A}(u) := \operatorname{supp}(u) = \{i \in [N] \colon u_i \neq 0\} \ .$$

For a fixed regularization parameter $t \geq 0$ and vector $f \in \mathbb{R}^m$, we define the set of minimizers $U_t(f)$ by

$$U_t(f) := \begin{cases} \operatorname{argmin}_{u\in\mathbb{R}^N} \frac{1}{2}\|Au - f\|_2^2 + t\ \|u\|_1 & \text{if } t > 0 \\ \operatorname{argmin}_{u\in\mathbb{R}^N} \|u\|_1 \quad \text{s.t. } A^T(Au - f) = 0 & \text{if } t = 0 \end{cases} \ . \tag{2}$$

We will often drop the dependence on f and simply write U_t.

We recall some basic facts about the variational problem (2). A proof is included for the reader's convenience.

Lemma 1 ([14])**.** *Let $u_1(t), u_2(t) \in U_t$ be two minimizers. Then one has that:*

(a) $Au_1(t) = Au_2(t)$;

(b) for $t > 0$, the map p given by

$$p(t) := \frac{1}{t}A^T(f - Au(t)) \tag{3}$$

satisfies $p(t) \in \partial\|u(t)\|_1$ *for all* $t > 0$ *and is independent of the specific choice of* $u(t) \in U_t$*;*

(c) $\|A^T(f - Au(t))\|_\infty \leq t \quad$ *and* $\quad \mathcal{A}(u(t)) \subseteq \mathcal{E}(t) := \mathcal{E}(t, u(t))$.

Remark 2. In the following, $p(t)$ always refers to the subgradient given by (3).

Proof. For $t = 0$, (a) holds since the constraint $A^TAu_1(t) = A^TAu_2(t)$ implies $u_1(t) - u_2(t) \in \mathrm{Ker}(A^TA) = \mathrm{Ker}(A)$.
For $t > 0$, (a) follows from the strict convexity of $\|\cdot\|_2^2$. Indeed, set $v := \frac{1}{2}u_1(t) + \frac{1}{2}u_2(t)$. Then

$$\begin{aligned} &\frac{1}{2}\|Av - f\|_2^2 + t\|v\|_1 \\ &\leq \frac{1}{2}\left(\frac{1}{2}\|Au_1(t) - f\|_2^2 + t\|u_1(t)\|_1\right) + \frac{1}{2}\left(\frac{1}{2}\|Au_2(t) - f\|_2^2 + t\|u_2(t)\|_1\right) \\ &= \min_{u\in\mathbb{R}^N}\left(\frac{1}{2}\|Au - f\|_2^2 + t\|u\|_1\right) . \end{aligned} \tag{4}$$

If $Au_1(t) \neq Au_2(t)$, the inequality (4) would be strict, leading to a contradiction. As a result, $Au(t)$ is independent of the minimizer chosen and $\mathcal{E}(t) = \mathcal{E}(t, u(t))$ is well-defined.

The statements (b) and (c) are an immediate consequence of the optimality condition

$$0 \in A^T(Au(t) - f) + t\ \partial\|u\|_1 ,$$

since the subdifferential of the ℓ_1-norm is given by

$$\partial\|u\|_1 = \{p \in \mathbb{R}^N \colon p_i = \operatorname{sgn} u_i\ \forall i \in \mathcal{A}(u) \text{ and } |p_i| \leq 1\ \forall i \notin \mathcal{A}(u)\} .$$

□

3 The Set of Possible Directions $\mathcal{D}$

We aim to construct a piecewise linear and continuous function $u \colon \mathbb{R}_{\geq 0} \to \mathbb{R}^N$ satisfying

$$u(t) \in U_t = \operatorname*{argmin}_{u\in\mathbb{R}^N} \frac{1}{2}\|Au - f\|_2^2 + t\ \|u\|_1 \quad \forall t > 0 \tag{5}$$

and $u(0) \in U_0$. To this end we make the following ansatz.
Assume we already have a solution $u(\hat{t}) \in U_{\hat{t}}$. Set $u(t) = u(\hat{t}) + (\hat{t} - t)\ d$ and try to choose $d \in \mathbb{R}^N$ such that $u(t) \in U_t$ for all $t \in [\hat{t} - \delta, \hat{t}]$, where $\delta = \delta(\hat{t}, u(\hat{t}), d) > 0$. This motivates the definition of the set of all possible directions $\mathcal{D}(\hat{t}, u(\hat{t}))$.

Definition 3. Let $\hat{t} > 0$ be a regularization parameter and $u(\hat{t}) \in U_{\hat{t}}$ be a solution of the variational problem (5). The set of all possible directions $\mathcal{D}(\hat{t}, u(\hat{t}))$ is defined as

$$\mathcal{D}(\hat{t}, u(\hat{t})) = \left\{d \in \mathbb{R}^N \colon \exists\delta \in (0, \hat{t}] \text{ s.t. } u(t) = u(\hat{t}) + (\hat{t} - t)d \in U_t\ \forall t \in [\hat{t} - \delta, \hat{t}]\right\} .$$

We now state and prove the main theorem of this section.

Theorem 4. *The set of possible directions* $\mathcal{D}(\hat{t}, u(\hat{t}))$ *at* $\left(\hat{t}, u(\hat{t})\right)$ *is the set of solutions to a non-negative least squares problem. More precisely, set*

$$r(\hat{t}) = f - Au(\hat{t}) \quad \textit{and} \quad p(\hat{t}) = \frac{1}{\hat{t}}A^T\left(f - Au(\hat{t})\right) .$$

Then we have that

$$\mathcal{D}(\hat{t}, u(\hat{t})) = \operatorname*{argmin}_{d\in\mathbb{R}^N} \|Ad - \frac{1}{\hat{t}} r(\hat{t})\|_2^2 \quad s.t. \quad d_i \; p(\hat{t})_i \geq 0 \;\forall i \in \mathcal{E}(\hat{t})\backslash\mathcal{A}(u(\hat{t})) \;, \tag{6}$$
$$d_i = 0 \;\forall i \notin \mathcal{E}(\hat{t}) \;.$$

Remark 5. The major difference to the standard homotopy method (cf. Algorithm 2) is the condition $d_i p(\hat{t})_i \geq 0$ for all $i \in \mathcal{E}(\hat{t})\backslash\mathcal{A}(u(\hat{t}))$. In fact, if $u(t) \in U_t$ for all $t \in [\hat{t}-\delta, \hat{t}]$, then the direction d necessarily satisfies this condition. To see this, let $i \in \mathcal{E}(\hat{t})\backslash\mathcal{A}(u(\hat{t}))$ and $p(\hat{t})_i = 1$. It follows that $u(t)_i \geq 0$ for all $t \in [\hat{t}-\delta, \hat{t}]$ because $p(t)$ is continuous and $p(\hat{t})_i = 1$. Therefore $u(\hat{t})_i + (\hat{t}-t)d_i \geq 0$, which, since $i \notin \mathcal{A}(u(\hat{t}))$, implies that $p(\hat{t})_i d_i = d_i \geq 0$.
As we will discuss in Section 5.1 below, this condition is sometimes violated for directions computed by the standard homotopy method; so an extra condition is indeed necessary.

Remark 6. By a change of variable, the constraints $p(\hat{t})_i d_i \geq 0$ are easily transformed into non-negativity constraints, which makes (6) a nonnegative least squares problem.

Proof. The strategy of the proof is as follows: We characterize the solutions of the nonnegative least squares problem by the Karush Kuhn Tucker (KKT) conditions (equations (7)-(13)), and compare them componentwise to a characterization of the set of possible directions.
Throughout the proof, let $\mathcal{E} := \mathcal{E}(\hat{t}),\ \mathcal{A} := \mathcal{A}(u(\hat{t})),$ and $\mathcal{D} := \mathcal{D}(\hat{t}, u(\hat{t}))$.

Let us start by stating the KKT conditions for the nonnegative least squares problem in (6): $d \in \mathbb{R}^N$ is a minimizer of (6) if and only if there exist $\lambda, \theta \in \mathbb{R}^N$ such that

$$A^T Ad - p(\hat{t}) + \lambda + \theta = 0 \;, \tag{7}$$
$$d_{\mathcal{E}^C} = 0 \;, \tag{8}$$
$$\theta_{\mathcal{E}} = 0 \;, \tag{9}$$
$$d_i \; p(\hat{t})_i \geq 0 \quad \forall i \in \mathcal{E}\backslash\mathcal{A} \;, \tag{10}$$
$$\lambda_{(\mathcal{E}\backslash\mathcal{A})^C} = 0 \;, \tag{11}$$
$$\lambda_i \; p(\hat{t})_i \leq 0 \quad \forall i \in \mathcal{E}\backslash\mathcal{A} \;, \tag{12}$$
$$\lambda_i \; d_i = 0 \quad \forall i \in \mathcal{E}\backslash\mathcal{A} \;. \tag{13}$$

We now show that every solution d of this system is a possible direction. We need to prove that there exists a $\delta > 0$ such that $u(t) = u(\hat{t}) + (\hat{t}-t)d \in U_t$ for all $t \in [\hat{t}-\delta, \hat{t}]$. Recalling the optimality condition

$$0 \in A^T(Au(t) - f) + t \; \partial\|u(t)\|_1 \;,$$

it suffices to show that $p(t) := \frac{1}{t} A^T(f - Au(t)) \in \partial\|u(t)\|_1$.
We begin by rewriting $p(t)$. By inserting the definition of $u(t)$, it follows that

$$\begin{aligned} p(t) &= \frac{1}{t} A^T(f - Au(\hat{t}) - (\hat{t}-t)Ad) \\ &= \frac{\hat{t}}{t} p(\hat{t}) - \frac{\hat{t}-t}{t} A^T Ad \\ &= p(\hat{t}) + \frac{\hat{t}-t}{t}(p(\hat{t}) - A^T Ad) \;. \end{aligned} \tag{14}$$

We argue componentwise proving that $p(t)_i \in \partial|u(t)_i|$ for all $i \in [N]$. We distinguish the three different cases $i \in \mathcal{A}$, $i \in \mathcal{E}\backslash\mathcal{A}$, and $i \in \mathcal{E}^C$.

Case 1: $i \in \mathcal{A}$. Since $u(t)$ is continuous, the equality $\operatorname{sgn}(u(t)_i) = \operatorname{sgn}(u(\hat{t})_i)$ holds for all t in some small interval $[\hat{t} - \delta_{\mathcal{A}}, \hat{t}]$. Using (7),(9), and (11), it follows that

$$p(t)_i = p(\hat{t})_i + \frac{\hat{t} - t}{t} \left(p(\hat{t}) - A^T A d\right)_i = p(\hat{t})_i \ .$$

Case 2: $i \in \mathcal{E} \backslash \mathcal{A}$. It follows that

$$u(t)_i = u(\hat{t})_i + (\hat{t} - t) d_i \overset{(10)}{=} (\hat{t} - t)|d_i| p(\hat{t})_i \ .$$

From (7), (9), and (14), we deduce that

$$p(t)_i = p(\hat{t})_i + \frac{\hat{t} - t}{t} \lambda_i \ .$$

If $d_i \neq 0$, then by complementary slackness (13), $\lambda_i = 0$ holds. Therefore $p(t)_i = p(\hat{t})_i$ and $\operatorname{sgn}(u(t)_i) = p(\hat{t})_i = p(t)_i$.
If $d_i = 0$, we have that $u(t)_i = 0$. Since $\lambda_i \, p(\hat{t})_i \leq 0$, it follows that for some $\delta_{\mathcal{E} \backslash \mathcal{A}} > 0$

$$|p(t)_i| = |p(\hat{t}) + \frac{\hat{t} - t}{t} \lambda_i| \leq 1 \quad \forall t \in [\hat{t} - \delta_{\mathcal{E} \backslash \mathcal{A}}, \hat{t}] \ .$$

Case 3: $i \in \mathcal{E}^C$. Lemma 1 and equation (8) yield $u(\hat{t})_{\mathcal{E}^C} = 0$ and $d_{\mathcal{E}^C} = 0$. Thus, it follows that $u(t)_i = 0$. Since $|p(\hat{t})_i| < 1$ and $p(t)$ is continuous, there exists a $\delta_{\mathcal{E}^C} > 0$ such that we have $|p(t)_i| \leq 1$ for all $t \in [\hat{t} - \delta_{\mathcal{E}^C}, \hat{t}]$.

Setting $\delta = \min\{\delta_{\mathcal{A}}, \delta_{\mathcal{E} \backslash \mathcal{A}}, \delta_{\mathcal{E}^C}\}$, we conclude that $p(t) \in \partial \|u(t)\|_1$ for all $t \in [\hat{t} - \delta, \hat{t}]$, and hence that d is a valid direction.

It remains to show that every possible direction d is a solution to the nonnegative least squares problem. To this end, we show that d satisfies the KKT conditions (7)-(13). Set

$$\begin{aligned} \lambda_{\mathcal{E} \backslash \mathcal{A}} &:= (p(\hat{t}) - A^T A d)_{\mathcal{E} \backslash \mathcal{A}} \ , \quad \lambda_{(\mathcal{E} \backslash \mathcal{A})^C} := 0 \ , \\ \theta_{\mathcal{E}^C} &:= (p(\hat{t}) - A^T A d)_{\mathcal{E}^C} \ , \qquad\qquad \theta_{\mathcal{E}} := 0 \ . \end{aligned}$$

Then (9) and (11) are satisfied by definition.
As there exists a $\delta > 0$ such that $u(t) \in U_t$ for all $[\hat{t} - \delta, \hat{t}]$, we conclude

$$p(t) = p(\hat{t}) + \frac{\hat{t} - t}{t}(p(\hat{t}) - A^T A d) = \frac{1}{t} A^T (f - A u(t)) \in \partial \|u(t)\|_1 \quad \forall t \in [\hat{t} - \delta, \hat{t}] \ . \tag{15}$$

Use this observation to first prove the multiplier equation (7), then the feasibility condition (8), and finally the equations (10), (12), and (13) concerning λ.

To prove (7), we need to show $(A^T A d)_{\mathcal{A}} = p(\hat{t})_{\mathcal{A}}$. Since $u(t)$ is continuous,

$$\operatorname{sgn}(u(t)_i) = \operatorname{sgn}(u(\hat{t})_i)$$

for all $t \in [\hat{t} - \delta, \hat{t}]$ and $i \in \mathcal{A}$. Therefore $p(t)_i = p(\hat{t})_i$, which together with (15) yields $(A^T A d)_i = p(\hat{t})_i$.
From Lemma 1 and the continuity of $p(t)$, it follows that $\mathcal{A}(u(t)) \subseteq \mathcal{E}(t) \subseteq \mathcal{E}(\hat{t}) = \mathcal{E}$. Thus

$\operatorname{supp}(d) \subseteq \mathcal{E}$, which proves (8).

To conclude, we prove (10), (12), and (13). For this, let $i \in \mathcal{E}\backslash\mathcal{A}$. First, assume that $d_i \neq 0$. Since $u(\hat{t})_i = 0$, it follows that

$$p(t)_i = \operatorname{sgn}(u(t)_i) = \operatorname{sgn}(d_i) \quad \forall t \in [\hat{t} - \delta, \hat{t}] \ .$$

Thus $p(t)_i$ is constant, and (15) yields $\lambda_i = (p(\hat{t}) - A^T A d)_i = 0$.
Second, assume that $d_i = 0$. Then (10) and (13) follow immediately, and

$$p(t)_i \overset{(15)}{=} p(\hat{t})_i + \frac{\hat{t} - t}{t}\lambda_i \ .$$

If (12) were violated, then for all $t \in [\hat{t} - \delta, \hat{t})$ we would have

$$|p(t)_i| = |p(\hat{t})_i + \frac{\hat{t} - t}{t}\lambda_i| > |p(\hat{t})_i| = 1 \ ,$$

which would contradict $p(t) \in \partial\|u(t)\|_1$. □

Remark 7. To implement the generalized homotopy method, we need an explicit expression for the maximal step size $\delta = \delta(\hat{t}, u(\hat{t}), d) > 0$. For this, define $s_{\mathcal{A}}, s_{\mathcal{E}\backslash\mathcal{A}}, s_{\mathcal{E}^C}, s \in \mathbb{R}$ as

$$\begin{aligned}
s_{\mathcal{A}} &:= \max_{i \in \mathcal{A}} \{\nu_i \colon d_i \neq 0 \text{ and } \nu_i < \hat{t}\} \text{ with } \nu_i := \frac{u(\hat{t})_i + \hat{t} d_i}{d_i} \ , \\
s_{\mathcal{E}\backslash\mathcal{A}} &:= \max_{i \in \mathcal{E}\backslash\mathcal{A}} \{\frac{|\lambda_i|}{|\lambda_i| + 2}\hat{t}\} \ , \\
s_{\mathcal{E}^C} &:= \max_{i \in \mathcal{E}^C} \{\mu_i^{\pm} \colon \mu_i^{\pm} < \hat{t}\} \text{ with } \mu_i^{\pm} := \begin{cases} \hat{t}\frac{(p(\hat{t}) - A^T A d)_i}{\pm 1 - (A^T A d)_i} & \text{if } \pm 1 \neq (A^T A d)_i \\ 0 & \text{else} \end{cases} \ , \\
s &:= \max\{s_{\mathcal{A}}, s_{\mathcal{E}\backslash\mathcal{A}}, s_{\mathcal{E}^C}, 0\} \ .
\end{aligned}$$

Then, the maximal step size is given by $\delta = \hat{t} - s$. This follows directly from the preceding proof.

Corollary 8. *There exist only finitely many sets of possible direction $\mathcal{D}(\hat{t}, u(\hat{t}))$.*

Proof. The corollary essentially follows from the KKT conditions (7)-(13) in the proof of Theorem 4. Recall that $d \in \mathcal{D}(\hat{t}, u(\hat{t}))$ is a possible direction if and only if there exist $(\lambda, \theta) \in \mathbb{R}^N \times \mathbb{R}^N$ such that (7)-(13) are satisfied. Since $\theta_{\mathcal{E}^C}$ can be chosen freely, the KKT conditions depend only on $\mathcal{E}, \mathcal{A},$ and $p(\hat{t})_{\mathcal{E}} \in \{\pm 1\}^{|\mathcal{E}|}$. Therefore, the set $\mathcal{D}(\hat{t}, u(\hat{t}))$ depends only on $\mathcal{E}, \mathcal{A},$ and $p(\hat{t})_{\mathcal{E}}$, which attain only finitely many different values. □

4 The Generalized Homotopy Method

The characterization of the set of possible directions $\mathcal{D}(\hat{t}, u(\hat{t}))$ in Theorem 4 directly yields a meta approach to compute a solution path: Start by choosing t^0 large enough to ensure that $u(t^0) = 0$ is a solution, i.e., $t^0 = \|A^T f\|_\infty$. Compute a direction $d^1 \in \mathcal{D}(t^0, u(t^0))$ and continue along the path $t \mapsto \big(t, u(t^0) + (t^0 - t)d^1\big) \in \mathbb{R} \times \mathbb{R}^N$ as long as $u(t) \in U_t$. Then, compute a new direction and repeat.
In the case of non-uniqueness, this approach yields a family of algorithms, as it needs to be combined with a rule R to choose a specific d from a given set $\mathcal{D}$ of potential directions, i.e., $d = R(\mathcal{D}) \in \mathcal{D}$.

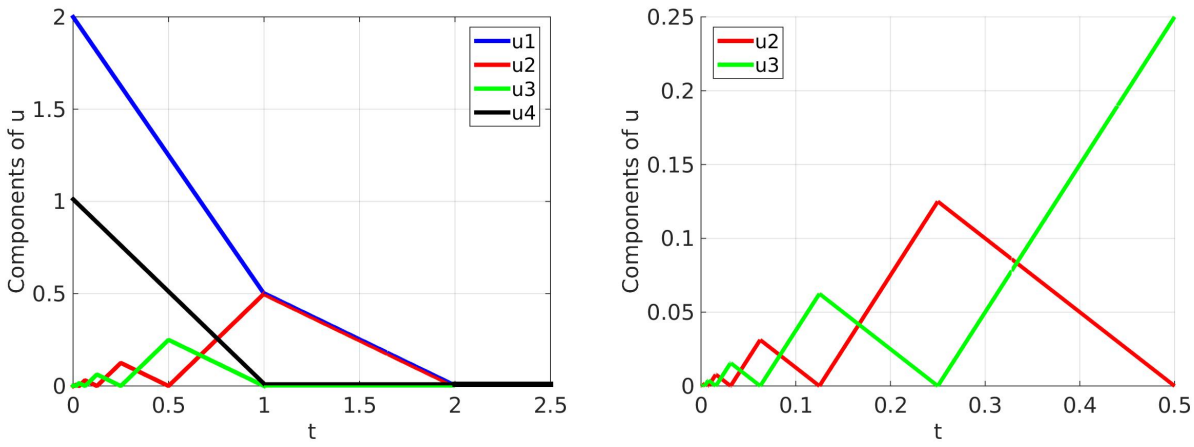

(a) Components of $u(t)$

(b) $u_2(t)$ and $u_3(t)$ on $[0, 1/2]$

Figure 2: For A and f as in (16), we display a piecewise linear and continuous solution path with infinitely many kinks.

The proof of the finite termination property [10, 6, 13] only holds for some and not for all of these algorithms as illustrated in Proposition 9 below. Thus for certain choice rules, the meta approach does not necessarily terminate after finitely many steps.

Proposition 9. *There exists a choice rule R, which, combined with the meta approach outlined above, yields a piecewise linear and continuous solution path $t \mapsto u(t)$ with infinitely many kinks for certain $A \in \mathbb{R}^{m\times N}$ and $f \in \mathbb{R}^m$.*

Proof. Let

$$A = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{2\times 4} \quad \text{and} \quad f = \begin{bmatrix} 2 \\ 1 \end{bmatrix} \in \mathbb{R}^2 \,. \tag{16}$$

Then $u(t) = (u_1(t), u_2(t), u_3(t), u_4(t))^T \in U_t$ if and only if

$$\begin{cases} u_1 = u_2 = u_3 = 0, & & u_4 = 0 & \text{if } t \geq 2 \\ u_1, u_2, u_3 \geq 0, & u_1 + u_2 + u_3 = 2 - t, & u_4 = 0 & \text{if } t \in (1, 2) \\ u_1, u_2, u_3 \geq 0, & u_1 + u_2 + u_3 = 2 - t, & u_4 = 1 - t & \text{if } t \in [0, 1] \end{cases} . \tag{17}$$

Thus already at the fist kink $t = 2$ there are multiple permissible directions, and it is not a priori clear which of them to choose. The choice $d = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 \end{bmatrix}^T$ is permissible, yielding $u_1(t) = u_2(t) = \frac{2-t}{2}$, $u_3(t) = u_4(t) = 0$. At $t = 1$, a change of direction is necessary to prevent that u_4 violates (17). A new permissible direction is $d = \begin{bmatrix} \frac{3}{2} & -1 & \frac{1}{2} & 1 \end{bmatrix}^T$. Now $u_2(t)$ decreases and hits 0 at $t = 2$, so again a change of direction is required; $d = \begin{bmatrix} \frac{3}{2} & \frac{1}{2} & -1 & 1 \end{bmatrix}^T$ is permissible. Continuing in this fashion and alternating between $d = \begin{bmatrix} \frac{3}{2} & -1 & \frac{1}{2} & 1 \end{bmatrix}^T$ and $d = \begin{bmatrix} \frac{3}{2} & \frac{1}{2} & -1 & 1 \end{bmatrix}^T$, one obtains kinks at $t = 2^{-k}$ for every $k \in \mathbb{N}$. It is easy to check that the three different directions chosen really correspond to different sets $\mathcal{D}$, so we are following a choice rule R.
The resulting solution path is displayed in Figure 2. □

To the best of our knowledge, this phenomenon was not discussed in any previous work dealing with the Lasso, nor have any specific choice rules been studied which avoid it. We propose to always choose the direction $d^{j+1} \in \mathcal{D}(t^j, u(t^j))$ with minimal ℓ_2-norm, which yields the generalized homotopy method (Algorithm 1). This approach is computationally feasible. For example, by first computing any $\tilde{d} \in \mathcal{D}(t^j, u(t^j))$, it can be formulated as

$$d^{j+1} \in \operatorname*{argmin}_{d \in \mathbb{R}^N} \|d\|_2^2 \quad \text{s.t. } Ad = A\tilde{d},\ d_{\mathcal{E}(t^j)^C} = 0, \\ d_i p(t^j)_i \geq 0 \ \forall i \in \mathcal{E}(t^j) \backslash \mathcal{A}(u(t^j)) \ . \tag{18}$$

In the most common scenarios (see Lemmas 17, 18 and 19), the computation of d^{j+1} is simpler than (18).

Algorithm 1 Generalized Homotopy Method

Input: data $f \in \mathbb{R}^m$, matrix $A \in \mathbb{R}^{m \times N}$
Output: number of steps K, sequence $t^0, \ldots, t^K$ of regularization parameters, sequence $u(t^0), \ldots, u(t^K)$ of solutions
Initialization: Set $t^0 = \|A^T f\|_\infty$ and $u(t^0) = 0$.
for $j = 0, 1, \ldots$ **do**
  if $t^j = 0$ **then**
    Break
  end if
  Compute $r^j = f - Au(t^j)$, $\mathcal{E} = \mathcal{E}(t^j)$, and $\mathcal{A} = \mathcal{A}(u(t^j))$.
  Set

$$\mathcal{D} = \operatorname*{argmin}_d \|Ad - \frac{1}{t^j} r^j\|_2^2 \quad \text{s.t. } d_{\mathcal{E}^C} = 0,\ d_i\ p(t^j)_i \geq 0 \ \forall i \in \mathcal{E} \backslash \mathcal{A} \ ,$$

  and compute $d^{j+1} = \operatorname{argmin}_{d \in \mathcal{D}} \|d\|_2^2$.
  Using Remark 7, find the minimal $t^{j+1} \geq 0$ s.t. $u(t) \in U_t$ for all $t \in [t^{j+1}, t^j]$.
end for

Let $t^0 = \|A^T f\|_\infty > t^1 > \ldots > t^K = 0$ and $u(t^0), u(t^1), \ldots, u(t^K)$ be the outputs of Algorithm 1. The path $u \colon \mathbb{R}_{\geq 0} \to \mathbb{R}^N,\ t \mapsto u(t)$ is then defined via linear interpolation by

$$u(t) = \begin{cases} 0 & \text{if } t \geq t^0 \\ \frac{t - t^k}{t^{k-1} - t^k} u(t^{k-1}) + \frac{t^{k-1} - t}{t^{k-1} - t^k} u(t^k) & \text{if } t \in [t^k, t^{k-1}) \end{cases} \ . \tag{19}$$

The following theorem, the main result of this paper, shows that $u(t)$ is indeed a solution path.

Theorem 10. *The generalized homotopy method (Algorithm 1) terminates after finitely many iterations. Furthermore, $u \colon \mathbb{R}_{\geq 0} \to \mathbb{R}^N$ as in (19) is piecewise linear, continuous, and satisfies*

$$u(t) \in U_t = \operatorname*{argmin}_{u \in \mathbb{R}^N} \frac{1}{2} \|Au - f\|_2^2 + t\ \|u\|_1$$

for all $t > 0$ as well as $u(0) \in U_0$.

To prove Theorem 10, we need the following lemma, which describes the dependence of the solution sets U_t on t.

Lemma 11. *Let $p(t)$ be the subgradient as in Lemma 1. For every $\mathcal{E} \subseteq [N]$ and every $s \in \{\pm 1\}^{|\mathcal{E}|}$ the set*

$$I_{\mathcal{E},s} = \{0 < t \leq \|A^T f\|_\infty \colon \mathcal{E} = \operatorname*{argmax}_{i \in [N]} |p(t)_i|,\ p(t)_\mathcal{E} = s\}$$

is an interval. If a, b with $a < b$ lie in the same $I_{\mathcal{E},s}$, and $u(a) \in U_a$ as well as $u(b) \in U_b$, then for every $t \in [a, b]$ the linear interpolation

$$u(t) = \frac{b-t}{b-a} u(a) + \frac{t-a}{b-a} u(b)$$

satisfies $u(t) \in U_t$.

Proof. Fix $\mathcal{E} \subseteq [N], s \in \{\pm 1\}^{|\mathcal{E}|}$, $a, b \in I_{\mathcal{E},s}$, and $u(a) \in U_a$ as well as $u(b) \in U_b$. Let $p(a)$ and $p(b)$ be the subgradients at $u(a)$ and $u(b)$ as defined in (3). Then in order to prove that $u(t) \in U_t$, we have to show that

$$\partial \|u(t)\|_1 \ni \tilde{p}(t) := \frac{1}{t} A^T (f - Au(t)) = \frac{b-t}{b-a} \frac{a}{t} p(a) + \frac{t-a}{b-a} \frac{b}{t} p(b)\ .$$

The last equality shows that $\tilde{p}(t)$ is a convex combination of $p(a)$ and $p(b)$ for every $t \in [a, b]$. Thus $\|\tilde{p}(t)_{\mathcal{E}^C}\|_\infty < 1$ and $\tilde{p}(t)_\mathcal{E} = p(a)_\mathcal{E} = p(b)_\mathcal{E} = s$. In particular, $\|\tilde{p}(t)\|_\infty \leq 1$. To show that $\tilde{p}(t) \in \partial \|u(t)\|_1$, it hence suffices to prove that $\|u(t)\|_1 \leq \langle \tilde{p}(t), u(t) \rangle$. Since $\mathcal{A}(u(a)) \cup \mathcal{A}(u(b)) \subseteq \mathcal{E}$, we have that

$$\begin{aligned}
\langle \tilde{p}(t), u(t) \rangle &= \frac{b-t}{b-a} \langle \tilde{p}(t), u(a) \rangle + \frac{t-a}{b-a} \langle \tilde{p}(t), u(b) \rangle \\
&= \frac{b-t}{b-a} \langle \tilde{p}(t)_\mathcal{E}, u(a)_\mathcal{E} \rangle + \frac{t-a}{b-a} \langle \tilde{p}(t)_\mathcal{E}, u(b)_\mathcal{E} \rangle \\
&= \frac{b-t}{b-a} \langle p(a)_\mathcal{E}, u(a)_\mathcal{E} \rangle + \frac{t-a}{b-a} \langle p(b)_\mathcal{E}, u(b)_\mathcal{E} \rangle \\
&= \frac{b-t}{b-a} \|u(a)\|_1 + \frac{t-a}{b-a} \|u(b)\|_1 \\
&\geq \|u(t)\|_1\ .
\end{aligned}$$

This concludes the proof of the second statement. In particular, $\tilde{p}(t)$ coincides with the subgradient $p(t)$ as in (3).
The first statement follows from $u(t) \in U_t$, $\|p(t)_{\mathcal{E}^C}\|_\infty < 1$, and $p(t)_\mathcal{E} = s$. □

Proof of Theorem 10. Recalling Theorem 4, d^{j+1} is indeed a feasible direction at $(t^j, u(t^j))$ provided that $u(t^j) \in U_{t^j}$. To see this, observe that

$$\{t \in [0, t^{j-1}] \colon u(t) = u(t^{j-1}) + (t^{j-1} - t) d^j \in U_t\}$$

is closed. Alternatively, one can also use explicit computation of the maximal $\delta > 0$ in Remark 7.

We show the finite termination property by contradiction, assuming that the algorithm does not terminate after finitely many steps. By Lemma 11 and the monotonicity of the t^j, there exists a set $\mathcal{E} \subseteq [N]$, a vector $s \in \{\pm 1\}^\mathcal{E}$, and a $j_0 \in \mathbb{N}$ such that $t^j \in I_{\mathcal{E},s}$ for all $j \geq j_0$.
We will now show that

$$\|d^{j+1}\|_2 < \|d^{j+2}\|_2 \qquad \forall\ j \geq j_0\ . \tag{20}$$

Since $t^j, t^{j+2} \in I_{\mathcal{E},s}$, there exists, again by Lemma 11, a $\bar{d} \in \mathcal{D}(t^j, u(t^j))$ such that

$$u(t^{j+2}) = u(t^j) + (t^j - t^{j+2})\bar{d} \ .$$

By construction,

$$u(t^{j+2}) = u(t^j) + (t^j - t^{j+1})d^{j+1} + (t^{j+1} - t^{j+2})d^{j+2} \ .$$

It follows that

$$\bar{d} = \frac{t^j - t^{j+1}}{t^j - t^{j+2}} d^{j+1} + \frac{t^{j+1} - t^{j+2}}{t^j - t^{j+2}} d^{j+2}$$

is a convex combination of d^{j+1} and d^{j+2}. If $\|d^{j+2}\|_2 \leq \|d^{j+1}\|_2$, we would have $\|\bar{d}\|_2 \leq \|d^{j+1}\|_2$ by convexity. But, since d^{j+1} is the unique direction with minimal ℓ_2-norm (by strict convexity), this yields $\bar{d} = d^{j+1}$, which is a contradiction to $d^{j+2} \neq d^{j+1}$. This shows (20).
By Corollary 8, there exist only finitely many sets of possible directions $\mathcal{D}(t^j, u(t^j))$. Since d^{j+1} is uniquely determined for each $\mathcal{D}(t^j, u(t^j))$, the set $\{d^{j+1} \colon j \in \mathbb{N}\}$ is also finite, a contradiction to (20). □

5 Relation to previous work

In this section, we compare the generalized homotopy method with previous homotopy algorithms [10, 6, 9, 13] and the adaptive inverse scale space method [2].

5.1 Standard Homotopy Method

At the core of te generalized homotopy method is a nonnegative least squares prolem to locally choose the direction. In contrast, previous works on the homotopy method [10, 6, 13] proposed to find a direction by solving a linear system. We will refer to the resulting algorithm as the standard homotopy method (Algorithm 2). Note that for reasons of better comparison to Algorithm 1, we have slightly extended the method to also accept inputs where the one-at-a-time condition (see Definition 12) does not hold. A first step towards dealing with scenarios where the one-at-a-time condition fails is the homotopy method with looping, which was, based on ideas in [6], introduced in [9], and is summarized in Algorithm 3.

Definition 12 ([6])**.** Let $t^0, \dots, t^K$ and $u(t^0), \dots, u(t^K)$ be the output produced by Algorithm 2. An index $i \in \mathrm{L}^j := \mathcal{A}(u(t^{j-1}))\backslash\mathcal{A}(u(t^j))$ is called a leaving coordinate, and an index $i \in \mathrm{H}^j := \mathcal{E}(t^j)\backslash\mathcal{E}(t^{j-1})$ is called a hitting coordinate.
We say the *one-at-a-time* condition is satisfied, if for every $0 \leq j \leq K$ such that $t^j > 0$ we have that

$$|\mathrm{H}^j \dot{\cup} \mathrm{L}^j| \leq 1 \ . \tag{22}$$

Theorem 13 ([6])**.** *Assume that the one-at-a-time condition holds and that $A_{\mathcal{E}(t^j)}$ is injective at every iteration. Then, the standard homotopy algorithm computes the unique solution path in finitely many steps.*

Our next result shows that under the same assumptions, the standard and generalized homotopy methods agree.

Theorem 14. *Assume that the one-at-a-time condition holds and that $A_{\mathcal{E}(t^j)}$ is injective at every iteration. Then the outputs of the standard homotopy method [10, 6] and the generalized homotopy method (Algorithm 1) coincide.*

Algorithm 2 Standard Homotopy Method [10, 6]

Input: data $f \in \mathbb{R}^m$, matrix $A \in \mathbb{R}^{m \times N}$
Output: number of steps K, sequence $t^0, \dots, t^K$ of regularization parameters, sequence $u(t^0), \dots, u(t^K)$ of solutions
Initialization: Set $t^0 = \|A^T f\|_\infty$ and $u(t^0) = 0$.
for $j = 0, 1, \dots$ **do**
  if $t^j = 0$ **then**
    Break
  end if
  Set $\mathrm{L}^j = \mathcal{A}(u(t^{j-1})) \backslash \mathcal{A}(u(t^j))$ and $\mathrm{S}^j = \mathcal{E}(t^j) \backslash \mathrm{L}^j$.
  Compute

$$d^{j+1}_{\mathrm{S}^j} = \left(A^T_{\mathrm{S}^j} A_{\mathrm{S}^j}\right)^\dagger p(t^j)_{\mathrm{S}^j} \quad \text{and set} \quad d^{j+1}_{(\mathrm{S}^j)^C} = 0 \ . \tag{21}$$

  Find the minimal $t^{j+1} \geq 0$ such that $u(t)$ solves (5) on $[t^{j+1}, t^j]$.
  if $t^{j+1} = t^j$ **then**
    Error "Algorithm failed to produce a solution path."
  end if
end for

Remark 15. As the injectivity assumption guarantees the uniqueness of the solution path, Theorem 14 directly follows from Theorem 10 and Theorem 13. Nevertheless, we provide a self-contained proof. This also yields an alternative proof of Theorem 13.

To prove Theorem 14, we first show Proposition 16, which states that d^{j+1} as in Algorithm 1 has the form (21) for some, a-priori unknown, set $\mathcal{S}^j$. The following three lemmas then show that the support set agrees with S^j as in Algorithm 2. They correspond to the three different cases in the proof of Theorem 4.

Proposition 16. *Let $0 \leq j \leq K-1$. Let $\mathcal{S}^j := \mathcal{A}(u(t^j)) \cup \mathrm{supp}(d^{j+1})$, where d^{j+1} is as in Algorithm 1. Then $\mathcal{A}(u(t^j)) \subseteq \mathcal{S}^j \subseteq \mathcal{E}(t^j) \cap \mathcal{E}(t^{j+1})$,*

$$d^{j+1}_{\mathcal{S}^j} = (A^T_{\mathcal{S}^j} A_{\mathcal{S}^j})^\dagger p(t^j)_{\mathcal{S}^j} \ , \quad \textit{and} \quad d^{j+1}_{(\mathcal{S}^j)^C} = 0 \ .$$

In light of Proposition 16, the standard homotopy method makes the educated guess $\mathcal{S}^j = \mathcal{E}(t^j) \backslash \mathrm{L}^j$, which is always correct under the assumptions in Theorem 13.

Proof. Throughout the proof, we write $\mathcal{E} = \mathcal{E}(t^j)$, $\mathcal{A} = \mathcal{A}(u(t^j))$, $\mathcal{S} = \mathcal{S}^j$ and $\mathcal{D} = \mathcal{D}(t^j, u(t^j))$. Set

$$\beta_{\mathcal{S}} = \left(A^T_{\mathcal{S}} A_{\mathcal{S}}\right)^\dagger p(t^j)_{\mathcal{S}} = \left(A^T_{\mathcal{S}} A_{\mathcal{S}}\right)^\dagger A^T_{\mathcal{S}} \frac{r^j}{t^j} \quad \text{and} \quad \beta_{\mathcal{S}^C} = 0.$$

We have to show $\beta = d^{j+1}$. Since the sign-constraint is only imposed on indices in $\mathcal{E} \backslash \mathcal{A}$, there exists an $\epsilon \in (0, 1)$ such that $\beta(\tau) = (1-\tau) d^{j+1} + \tau \beta$ is in the feasible set of the nonnegative least squares problem (6). Further, by the KKT conditions (7), (9), and (13), we have

$$A^T_{\mathcal{S}} A_{\mathcal{S}} d^{j+1}_{\mathcal{S}} = p(t^j)_{\mathcal{S}} = A^T_{\mathcal{S}} A_{\mathcal{S}} \beta_{\mathcal{S}} \ .$$

Thus $A d^{j+1} = A\beta$ holds, and $\beta(\tau) \in \mathcal{D}$. By the definition of β, $\|\beta\|_2 \leq \|d^{j+1}\|_2$ holds, which yields $\|\beta(\tau)\|_2 \leq \|d^{j+1}\|_2$ for all $\tau \in [0, \epsilon]$. Since $\beta(\tau) \in \mathcal{D}$, this yields $\beta(\tau) = d^{j+1}$. Recalling the definition

of $\beta(\tau)$, the equality $\beta = d^{j+1}$ follows.

The inclusion $\mathcal{A}(u(t^j)) \subset \mathcal{S}^j$ holds by the definition of $\mathcal{S}^j$, and $\mathcal{S}^j \subseteq \mathcal{E}(t^j)$ folows from Theorem 4. To see that $\mathcal{S}^j \subseteq \mathcal{E}(t^{j+1})$, we distinguish two cases. If $t^{j+1} > 0$, then

$$p(t^{j+1})_{\mathcal{S}^j} \overset{(14)}{=} p(t^j)_{\mathcal{S}^j} + \frac{t^j - t^{j+1}}{t^{j+1}}(p(t^j)_{\mathcal{S}^j} - A^T_{\mathcal{S}^j} A d^{j+1}) = p(t^j)_{\mathcal{S}^j} \ ,$$

showing that $\mathcal{S}^j \subseteq \mathcal{E}(t^{j+1})$. Since $\mathcal{E}(0) = [N]$, $\mathcal{S}^j \subseteq \mathcal{E}(t^{j+1})$ also holds if $t^{j+1} = 0$. □

The first lemma describes a case of a leaving coordinate , i.e., that a coordinate $u(t)_i$ becomes zero which was previously non-zero.

Lemma 17. *Let* $\mathrm{H}^j = \emptyset$, $\mathrm{L}^j = \{i\}$, $\mathcal{E}(t^j) = \mathcal{A}(u(t^j))\dot\cup\{i\}$ *and* d^{j+1} *as in Algorithm 1. Then one has that*

$$d^{j+1}_{\mathcal{E}(t^j)\setminus\{i\}} = \left(A^T_{\mathcal{E}(t^j)\setminus\{i\}} A_{\mathcal{E}(t^j)\setminus\{i\}}\right)^\dagger p(t^j)_{\mathcal{E}(t^j)\setminus\{i\}} \quad \textit{and} \quad d^{j+1}_{(\mathcal{E}(t^j)\setminus\{i\})^C} = 0 \ .$$

Furthermore, if $A_{\mathcal{E}(t^j)}$ *is injective, then* $A_i^T A d^{j+1} \neq p(t^j)_i$ *and the index* i *leaves the equicorrelation set* $\mathcal{E}(t)$.

Proof. We set $\mathcal{S}^j$ as in Proposition 16. Since $\mathcal{E}(t^j) = \mathcal{A}(u(t^j))\dot\cup\{i\}$, either $\mathcal{S}^j = \mathcal{A}(u(t^j))$ or $\mathcal{S}^j = \mathcal{E}(t^j)$. From $u(t^{j-1})_i \neq 0$ (by the definition of L^j) and $u(t^j)_l \neq 0$ for all $l \in \mathcal{E}(t^j)\setminus\{i\} = \mathcal{A}(u(t^j))$ together with Proposition 16, it follows that $\mathcal{S}^{j-1} = \mathcal{E}(t^j)$. Now $\mathcal{S}^j \neq \mathcal{E}(t^j) = \mathcal{S}^{j-1}$, as otherwise, again by Proposition 16, $d^{j+1} = d^j$, which is a contradiction.
Now assume that $A_{\mathcal{E}(t^j)}$, and hence also $A^T_{\mathcal{E}(t^j)} A_{\mathcal{E}(t^j)}$, is injective and that $A_i^T A d^{j+1} = p(t^j)_i$. Then

$$A^T_{\mathcal{E}(t^j)} A_{\mathcal{E}(t^j)} d^{j+1}_{\mathcal{E}(t^j)} = p(t^j)_{\mathcal{E}(t^j)} = A^T_{\mathcal{E}(t^j)} A_{\mathcal{E}(t^j)} d^j_{\mathcal{E}(t^j)} \ ,$$

and thus $d^{j+1} = d^j$, which is again a contradiction. □

The second lemma deals with the case that $\mathrm{H}^j = \emptyset$ and $\mathrm{L}^j = \emptyset$. Under the additional assumption that $\mathcal{E}(t^j) = \mathcal{A}(u(t^j))\dot\cup\{i\}$, this implies that $i \in [N]$ is in the equicorrelation set for both $t = t^{j-1}$ and $t = t^j$, but on the interval $[t^j, t^{j-1}]$ the i-th component $p(t)_i$ changes from $+1$ to -1 or vice versa while $u(t)_i$ remains zero.

Lemma 18. *Assume that* $\mathcal{E}(t^j) = \mathcal{E}(t^{j-1})$ *and that* $i \in \mathcal{E}(t^j)$ *is an index such that* $\mathcal{A}(u(t^j)) = \mathcal{A}(u(t^{j-1})) = \mathcal{E}(t^j)\setminus\{i\}$ *and* $p(t^j)_i \neq p(t^{j-1})_i$. *Then*

$$d^{j+1}_{\mathcal{E}(t^j)} = \left(A^T_{\mathcal{E}(t^j)} A_{\mathcal{E}(t^j)}\right)^\dagger p(t^j)_{\mathcal{E}(t^j)} \quad \textit{and} \quad d^{j+1}_{(\mathcal{E}(t^j))^C} = 0 \ .$$

Furthermore, $d_i^{j+1} p(t^j)_i > 0$.

Proof. Again, for $\mathcal{S}^j$ as in Proposition 16, one has that either $\mathcal{S}^j = \mathcal{E}(t^j)\setminus\{i\}$ or $\mathcal{S}^j = \mathcal{E}(t^j)$. If $\mathcal{S}^{j-1} = \mathcal{E}(t^{j-1})$, Proposition 16 would yield

$$A^T_{\mathcal{E}(t^{j-1})} A d = p(t^{j-1})_{\mathcal{E}(t^{j-1})} \ ,$$

implying that $p(t^{j-1})_i = (A^T A d^j)_i$, and hence, with (14), a contradiction to the assumption $p(t^{j+1})_i \neq p(t^j)_i$. Thus, $\mathcal{S}^{j-1} = \mathcal{E}(t^j)\setminus\{i\}$. Since $d^{j+1} \neq d^j$ and $p(t^j)_{\mathcal{A}(u(t^{j-1}))} = p(t^{j-1})_{\mathcal{A}(u(t^{j-1}))}$, it follows that $\mathcal{S}^j \neq \mathcal{S}^{j-1}$, i.e., $\mathcal{S}^j = \mathcal{E}(t^j)$, which proves the first part of the lemma.
Lastly, since $d_i^{j+1} p(t^j)_i \geq 0$ by Algorithm 1 and $i \in \mathcal{S}^j\setminus\mathcal{A}(u(t^j)) \subseteq \mathrm{supp}(d^{j+1})$, it follows that $d_i^{j+1} p(t^j)_i > 0$. □

The third lemma describes the case of a hitting coordinate, i.e., a coordinate which has to be included in the equicorrelation set at $t = t^j$.

Lemma 19. *Let* $\mathrm{H}^j = \{i\},\ \mathrm{L}^j = \emptyset$, *and* $\mathcal{E}(t^j) = \mathcal{A}(u(t^j))\dot\cup\{i\}$. *Then we have that*

$$d^{j+1}_{\mathcal{E}(t^j)} = \left(A^T_{\mathcal{E}(t^j)} A_{\mathcal{E}(t^j)}\right)^\dagger p(t^j)_{\mathcal{E}(t^j)}\ , \quad d^{j+1}_{\mathcal{E}(t^j)^C} = 0\ ,$$

and $d^{j+1}_i p(t^j)_i > 0$.

Proof. By assumption $\mathcal{E}(t^j) = \mathcal{A}(u(t^j))\dot\cup\{i\}$. Then, either $\mathcal{S}^j = \mathcal{A}(u(t^j))$ or $\mathcal{S}^j = \mathcal{E}(t^j)$. To prove the lemma, it suffices to show that $\mathcal{S}^j \neq \mathcal{A}(u(t^j))$.
Since $u(t^j)_l \neq 0$ for all $l \in \mathcal{A}(u(t^j))$ and $\mathrm{L}^j = \emptyset$, it follows that $\mathcal{S}^{j-1} = \mathcal{A}(u(t^j))$. If $\mathcal{S}^j = \mathcal{A}(u(t^j))$, then $d^{j+1} = d^j$, which is a contradiction. The last inequality follows as in Lemma 18. □

Lemma 19 was already proven in [6, Lemma 5.4] under the additional assumption that $A_{\mathcal{E}(t^j)}$ is injective. The most difficult part in their proof is to show that d^{j+1}_i agrees in sign with $p(t^j)_i$, i.e., $d^{j+1}_i p(t^j)_i \geq 0$.

Proof of Theorem 14. The result follows from the first parts of the Lemmas 17, 18 and 19 if we show that the assumption $|\mathcal{E}(t^j)\backslash\mathcal{A}(u(t^j))| = 1$ is satisfied for every $j = 0, \ldots, K-1$. Due to the one-at-a-time condition, it suffices to note that $|\mathcal{E}(t)\backslash\mathcal{A}(u(t))| = 0$ for every $t \in (t^{j+1}, t^j)$, which follows directly from the second parts of the Lemmas 17, 18 and 19. □

Remark 20. Our analysis shows that the injectivity of $A_{\mathcal{E}(t^j)}$ is only needed to show that $1 = |\mathcal{E}(t^j)\backslash\mathcal{A}(u(t^j))|$ at every iteration, and only in the scenario of Lemma 17. Essentially, we have to exclude that a leaving index $i \in \mathrm{L}^j = \mathcal{A}(u(t^{j-1}))\backslash\mathcal{A}(u(t^j))$ remains in the equicorrelation set, i.e., $i \in \mathcal{E}(t)$ for all $t \in (t^{j+1}, t^j)$. As long as the solution path is unique, this would contradict Lemma 11. In the case of non-uniqueness, there may be additional kinks in the interior of one of the $I_{\mathcal{E},s}$ (as defined in Lemma 11), where we do not know yet whether and how it can be excluded.

It was noted in [6, 9] that without the one-at-a-time condition, the standard homotopy method can encounter sign inconsistencies. An example with $A \in \mathbb{R}^{3\times 3}$ is given in [9], namely

$$A = \begin{bmatrix} -3 & 4 & 4 \\ -5 & 1 & 4 \\ 5 & 1 & -4 \end{bmatrix} \quad \text{and} \quad f = \begin{bmatrix} 24 \\ 17 \\ -7 \end{bmatrix}\ . \tag{23}$$

The outputs of the generalized homotopy method and SparseLab [4], one of the most popular implementations of the homotopy method [10, 6], are displayed in Figure 3. At the first kink $t = 192$, the 3^{rd} component in the output produced by SparseLab enters the active set with the wrong sign. Therefore, in this example SparseLab is unable to produce a full solution path. As Algorithm 2 has an additional feature to detect sign inconsistencies, it would exit at $t = 192$.
Notice that the matrix A is invertible, and thus the solution path is unique. The wrong solution path produced by SparseLab is solely the result of the missing one-at-a-time condition, and not of non-uniqueness.

In [9], based on ideas in [6], the following strategy was proposed: Instead of choosing $\mathrm{S}^j = \mathcal{E}(t^j)\backslash\mathrm{L}^j$ in Algorithm 2, loop over all sets $\mathrm{S} \subseteq [N]$ with $\mathcal{A}(u(t^j)) \subseteq \mathrm{S} \subseteq \mathcal{E}(t^j)$, compute

$$d_{\mathrm{S}} = \left(A^T_{\mathrm{S}} A_{\mathrm{S}}\right)^\dagger p(t^j)_{\mathrm{S}}, \quad \text{and set} \quad d_{\mathrm{S}^C} = 0\ .$$

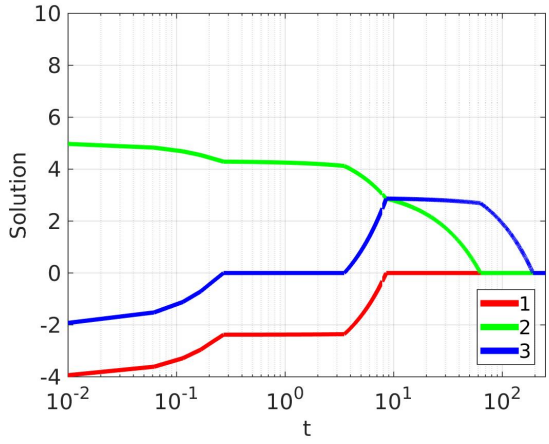

(a) Output of Algorithm 1

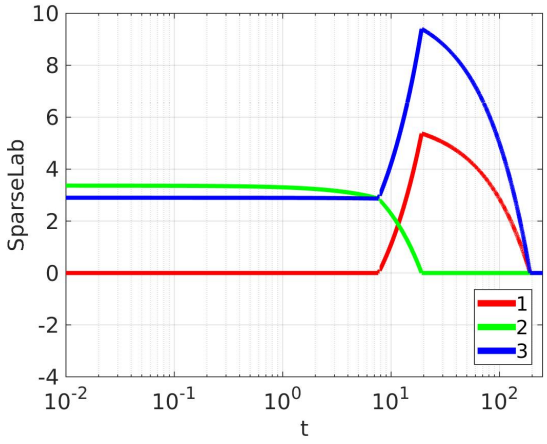

(b) Output of SparseLab

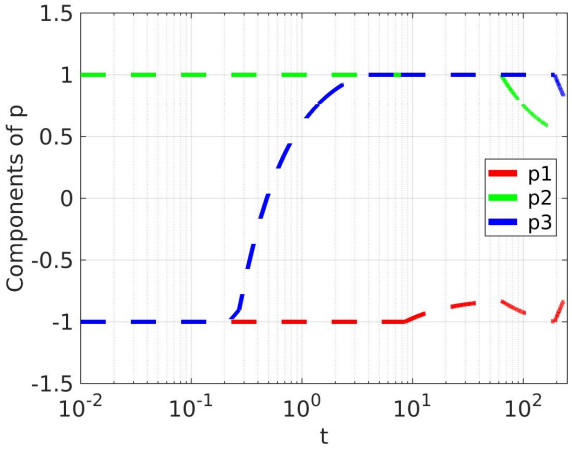

(c) Components of $p(t)$

Figure 3: We display the behaviour of the generalized homotopy method and the output of SparseLab in an example where the one-at-a-time condition does not hold. On certain intervals, the output of SparseLab has a non-zero first component while the correct subgradient satisfies $|p_1(t)| < 1$. Thus, SparseLab does not produce a solution path.

Choose $\mathrm{S}^j = \mathrm{S}$ as soon as $t^{j+1} < t^j$, i.e., $d \in \mathcal{D}(t^j, u(t^j))$. We call this the homotopy method with looping (see Algorithm 3).

From Proposition 16 it follows that the homotopy method with looping finds a direction at every iteration. This is, at least in the case of non-uniqueness, a non-trivial result. Notice that the homotopy method with looping does not necessarily compute the direction with minimal ℓ_2 norm. Therefore, the finite termination of the algorithm is unclear.
Besides providing a theoretical foundation to the homotopy method with looping, the characterization of the set of possible directions (Theorem 4) can also improve its performance. The loop over all sets $\mathrm{S} \subseteq [N]$ with $\mathcal{A}(u(t^j)) \subseteq \mathrm{S} \subseteq \mathcal{E}(t^j)$ can be interpreted as a rudimentary active-set strategy to solve the nonnegative least squares problem (6), even though this was not explicitly noted. As soon as $|\mathcal{E}(t^j) \backslash \mathcal{A}(u(t^j))|$ becomes large this methods becomes infeasible. Indeed, we would have to solve $2^{|\mathcal{E}(t^j) \backslash \mathcal{A}(u(t^j))|}$ linear systems. Empirical tests show that small random Bernoulli matrices, for

Algorithm 3 Homotopy Method With Looping [6, 9]

Input: data $f \in \mathbb{R}^m$, matrix $A \in \mathbb{R}^{m \times N}$
Output: number of steps K, sequence $t^0, \dots, t^K$ of regularization parameters, sequence $u(t^0), \dots, u(t^K)$ of solutions
Initialization: Set $t^0 = \|A^T f\|_\infty$ and $u(t^0) = 0$.
for $j = 0, 1, \dots$ **do**
  if $t^j = 0$ **then**
    Break
  end if
  for $S \subseteq [N]$ with $\mathcal{A}(u(t^j)) \subseteq S \subseteq \mathcal{E}(u(t^j))$ **do**
    Compute

$$d_{\mathrm{S}}^{j+1} = \left(A_{\mathrm{S}}^T A_{\mathrm{S}}\right)^\dagger p(t^j)_{\mathrm{S}} \quad \text{and set} \quad d_{(\mathrm{S})^C}^{j+1} = 0 \ .$$

    Find the minimal $t^{j+1} \geq 0$ such that $u(t)$ solves (5) on $[t^{j+1}, t^j]$.
    if $t^{j+1} < t^j$ **then**
      Break
    end if
  end for
end for

instance $A \in \mathbb{R}^{20 \times 50}$, regularly yield $|\mathcal{E}(t^j) \backslash \mathcal{A}(u(t^j))| \geq 18$. Nevertheless we consider their work [9] an important step towards understanding the solution paths of (5) even when the one-at-a-time condition fails.

The injectivity assumption in Theorem 13 is mainly needed to prevent the non-uniqueness of the solution path. Scenarios without uniqueness assumptions were studied in [13], but again only under the (implicit) assumption of the one-at-a-time condition. The results [13, Lemma 9 and Section 3.1] state that a continuous and piecewise linear solution path is given by the semi-explicit formula

$$\beta_{\mathcal{E}(t)} = (A_{\mathcal{E}(t)})^\dagger \left(f - (A_{\mathcal{E}(t)}^T)^\dagger \ t \ p(t)_{\mathcal{E}(t)}\right) \quad \text{and} \quad \beta_{(\mathcal{E}(t))^C} = 0 \ . \tag{24}$$

Although $\beta(t) \in U_t$ for A and f as in the previous example (23), this approach in general only applies under the one-at-a-time condition. As an example, consider

$$A = \begin{bmatrix} -1 & +1 & +1 & +1 \\ +1 & -1 & +1 & +1 \\ +1 & +1 & +1 & -1 \end{bmatrix}, \qquad f = \begin{bmatrix} -1 \\ -3 \\ -1 \end{bmatrix} \quad \text{and} \quad t = 2 \ . \tag{25}$$

Then $u(2) = \begin{bmatrix} 0 & 0 & -1 & 0 \end{bmatrix}^T$ is a solution and $p(2) = \begin{bmatrix} -1 & 1 & -1 & -1 \end{bmatrix}^T$ is the corresponding subgradient. But (24) yields $\beta(2) = \begin{bmatrix} -1/4 & -1/4 & -3/4 & -1/4 \end{bmatrix}$, and thus the second component has the wrong sign. As can be seen in Figure 4, the path $\beta(t)$ is neither continuous nor does it solve $\beta(t) \in U_t$ for all $t \geq 0$.

5.2 Adaptive Inverse Scale Space Method

The adaptive inverse scale space (aISS) method is a fast algorithm to compute ℓ_1-minimizing solutions of linear systems. Instead of calculating the minimizers of variational problems with varying

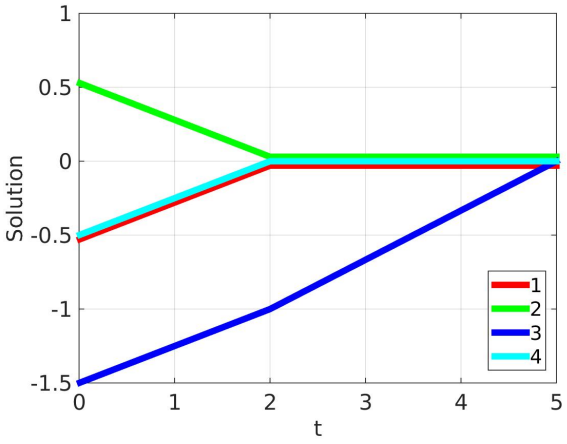

(a) Components of $u(t)$

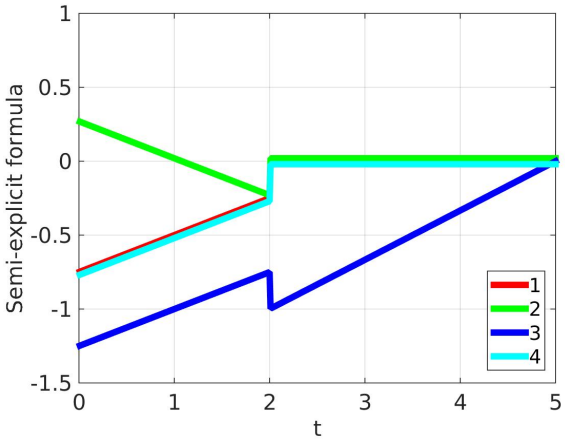

(b) Components of $\beta(t)$

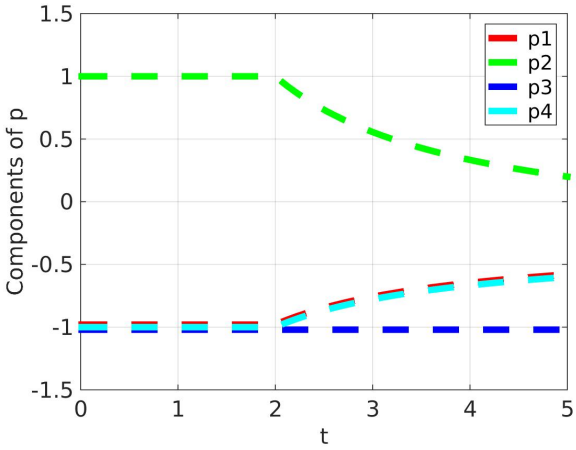

(c) Components of $p(t)$

Figure 4: For A and f as in (25) we compare a solution path $u(t)$ to the semi-explicit $\beta(t)$ as in (24). In this example, $\beta(t)$ is not a solution path.

regularization parameters t, it computes an exact solution to the differential inclusion

$$\begin{cases} \partial_\nu q(t) = A^T(f - Av(t)) \quad \text{with} \quad q(t) \in \partial\|v(t)\|_1 \\ q(0) = 0 \end{cases} . \tag{26}$$

The following theorem is proven in [2, Theorem 1 and 2].

Theorem 21 ([2]). *There exists a finite sequence of times*

$$0 = t^0 < t^1 < t^2 < \ldots < t^K < t^{K+1} = \infty$$

such that for all $k = 0, \ldots, K$

$$v(t) = v(t^k), \quad q(t) = q(t^k) + (t - t^k)\ A^T(f - Av(t^k)),$$

for $t \in [t^k, t^{k+1})$ *is a solution of the inverse scale space flow. Here,* $v(t^k)$ *is a solution of*

$$v(t^k) \in \operatorname*{argmin}_{v \in \mathbb{R}^N} \|Av - f\|_2^2 \quad \text{s.t. } q(t^k) \in \partial\|v\|_1 \ . \tag{27}$$

Furthermore, $v(t)$ is an ℓ_1-minimizing solution of $A^TAv = A^Tf$ for all $t \geq t^K$.

The aISS method has striking similarities to the generalized homotopy method. First, a seemingly continuous problem, i.e., a differential inclusion, can be solved completely by knowing the solution at finitely many points. While the generalized homotopy method produces a piecewise linear path $u(t)$, the path $v(t)$ of the aISS method is piecewise constant.
Second, both methods solve nonnegative least squares problems to calculate the solution path. To see this, note that in the aISS method

$$q(t^k) \in \partial\|v\|_1 \;\Leftrightarrow\; v_i q(t^k)_i \geq 0 \;\forall i \text{ with } |q(t^k)_i| = 1,\; v_i = 0 \;\forall i \text{ with } |q(t^k)_i| < 1 \;.$$

Although in a different context, the link between the inverse scale space flow and variational methods is also studied in [1].

6 Conclusions and Future Research

In this paper, we have introduced a generalized homotopy method which computes a full solution path of ℓ_1-regularized problems in finitely many iterations. In contrast to previous homotopy methods, it provably works for an arbitrary combination of a measurement matrix and a data vector, requiring neither the uniqueness of the solution path nor the one-at-a-time condition. The backbone of the generalized homotopy method is a characterization of the set of possible directions by a nonnegative least squares problem.
In future research, we will extend the proposed homotopy method to arbitrary polyhedral regularizations. Furthermore, we will investigate its applicability for generalizing the ideas of nonlinear spectral decompositions considered in [8, 1] to more general data fidelity terms.

Acknowledgements

DC and MM were supported by the ERC Starting Grant "ConvexVision". FK's contribution was supported by the German Science Foundation DFG in context of the Emmy Noether junior research group KR 4512/1-1 (RaSenQuaSI).

References

[1] M. Burger, G. Gilboa, M. Moeller, L. Eckardt, and D. Cremers, *Spectral Decompositions using One-Homogeneous Functionals*, (2016), pp. 1–31, arXiv:1601.02912.

[2] M. Burger, M. Möller, M. Benning, and S. Osher, *An Adaptive Inverse Scale Space Method for Compressed Sensing*, Mathematics of Computation, 82 (2013), pp. 269–299.

[3] E. Candes, J. Romberg, and T. Tao, *Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information*, IEEE Transactions on Information Theory, 52 (2006), pp. 489–509.

[4] D. Donoho, I. Drori, V. Stodden, Y. Tsaig, and M. Shahram, *Sparselab*. https://sparselab.stanford.edu/, 2007.

[5] D. L. Donoho and Y. Tsaig, *Fast Solution of ℓ_1-Norm Minimization Problems When the Solution May Be Sparse*, IEEE Transactions on Information Theory, 54 (2008), pp. 4789–4812.

[6] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, *Least Angle Regression*, The Annals of Statistics, 32 (2004), pp. 407–499.

[7] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, Birkhäuser, 2013.

[8] G. Gilboa, *A Total Variation Spectral Framework for Scale and Texture Analysis*, SIAM Journal on Imaging Sciences, 7 (2014), pp. 1937–1961.

[9] I. Loris, *L1Packv2: A Mathematica package for minimizing an ℓ_1-penalized functional*, Computer Physics Communications, 179 (2008), pp. 895–902.

[10] M. R. Osborne, B. Presnell, and B. A. Turlach, *A New Approach to Variable Selection in Least Squares Problems*, IMA Journal of Numerical Analysis, 20 (2000), pp. 389–403.

[11] L. I. Rudin, S. Osher, and E. Fatemi, *Nonlinear total variation based noise removal algorithms*, Physica D: Nonlinear Phenomena, 60 (1992), pp. 259–268.

[12] R. Tibshirani, *Regression Shrinkage and Selection Via the Lasso*, Journal of Royal Statistical Society, Series B, 58 (1996), pp. 267–288.

[13] R. J. Tibshirani, *The lasso problem and uniqueness*, Electronic Journal of Statistics, 7 (2013), pp. 1456–1490.

[14] H. Zhang, M. Yan, and W. Yin, *One condition for solution uniqueness and robustness of both l1-synthesis and l1-analysis minimizations*, (2013), pp. 1–15, arXiv:1304.5038.

[15] H. Zhang, W. Yin, and L. Cheng, *Necessary and Sufficient Conditions of Solution Uniqueness in 1-Norm Minimization*, Journal of Optimization Theory and Applications, 164 (2015), pp. 109–122.